

The Sensor Web: A New Instrument Concept

Kevin A. Delin* and Shannon P. Jackson

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109-8099

ABSTRACT

The confluence of the rapidly expanding sensor, computation, and telecommunication industries has allowed for a new instrument concept: the Sensor Web. A Sensor Web consists of intra-communicating, spatially-distributed sensor pods that are deployed to monitor and explore environments. It is capable of automated reasoning for it can perform intelligent autonomous operations in uncertain environments, respond to changing environmental conditions, and carry out automated diagnosis and recovery. Sensor Webs could have as much an impact on the uses of sensors as the Internet did on the uses of computers.

Sensor Webs are often confused with “distributed sensors” or “sensor networks”. The unique feature of the Sensor Web is that information gathered by one pod is shared and used by other pods. In contrast, sensor networks merely gather data and information gathered by a particular pod on such a network does not influence the behavior of another pod. Thus, sensor networks collect data while Sensor Webs can react and modify their behavior on the basis of the collected data.

This paper will outline the potential of the Sensor Web concept and describe the Jet Propulsion Laboratory Sensor Webs Project (<http://sensorwebs.jpl.nasa.gov/>). In particular, a prototype Sensor Web deployed at the Huntington Botanical Gardens will be discussed.

Keywords: Sensor Web, Distributed, Network, Smart Sensors, Wireless, Environment Monitoring, Botanical

1. INTRODUCTION

The present rapid developments in computation and telecommunication technologies have revolutionized the way in which we think about hardware. The requirements for today's computers and cellular telephones are so complex that the only possible way to develop these technologies at all is to sell them to a mass market and rely on the economy of scale to provide the necessary financial capital for the next round of development. As a result, we are living in an era where mass-produced, commercially available components often represent the state-of-the-art. This is in marked contrast to previous times where the state-of-the-art was represented almost exclusively in government, military, or university laboratories. Today, for example, hardware is so inexpensive that telephone companies will deeply discount, or even give away for free, cellular phones in anticipation of recouping the costs via telephone service.

The Sensor Web is a new type of instrument that exploits the availability of this low-cost, but extremely advanced, hardware for its platform. In fact, the Sensor Web builds on the present revolutions of computation, telecommunication, and sensing technologies. Additionally, the Sensor Web provides a new means to think about environmental monitoring and creating a virtual presence. For this reason, the Sensor Web is a new instrument concept, capable of being developed into a wide range of applications. This paper will first carefully describe what a Sensor Web is and how it is distinct from seemingly similar constructions such as spatially distributed sensors. A brief introduction to Sensor Web applications will follow with emphasis to those of interest to NASA. Finally, Sensor Webs we have developed and deployed will be described in detail.

2. WHAT IS A SENSOR WEB?

As defined in the NASA New Technology Report on Sensor Webs,¹ the Sensor Web consists of a system of wireless, intra-communicating, spatially distributed sensor pods that can be easily deployed to monitor and explore new environments (see Figure 1). Each pod consists of two primary modules. The first module comprises the transducers that physically interact with the environment and convert environmental parameters into electrical signals. The second module represents the infrastructure of the Sensor Web itself. Included in this module are telecommunication capabilities, power sources and

* To whom correspondence should be addressed: Mail stop: M/S 306-336, Email: Kevin.Delin@jpl.nasa.gov.

energy harvesting devices, and computation devices to run the protocol schemes and provide for local data analysis. The penultimate goal of a Sensor Web is to extract *knowledge* from the data collected and adapt and react accordingly. Although the computation hardware in a pod can be quite sophisticated, it is the sharing of information among the pods that gives the Sensor Web a macrointelligence. Intelligence in the human brain is created because of a complex, inhomogeneous network of neurons² and not because of individual intelligence from each neuron. Similarly, the Sensor Web is an instrument where greater functionality is derived from a parallel-type architecture as sensor measurements (including pod location) are passed, *and collectively interpreted*, from pod to pod. This global sharing of information will lead to pod synergism (the whole of their activity being greater than the sum of their parts) by permitting intelligent resource (power, bandwidth, consumables) management by the web, and allowing for self-modifying behavior based on environmental factors and internal web diagnostics.

The wireless communication between pods is presumed to be omni-directional. Unlike star-network configurations where data collected from all pods is passed directly to a central point, information within the Sensor Web is passed to an uplink point, denoted as a prime or mother pod, by hopping it from pod to pod. In other words, data from various pods *are shared* as well as communicated throughout the entire web. The overall protocol is quite simple. Information is obtained at each pod via two routes: (a) direct measurements taken by local sensors at that pod and (b) information gathered by other pods and communicated throughout the web. The key concept is that there is no artificial differentiation between the two types of information. The protocol then is to simply rebroadcast the data (actual measurements) or information (digested data) to any pod within communication range. Any information received at a mother pod is not rebroadcast to the daughter pods and disappears from the web at this point although it is accessible to an outside user or another mother pod.

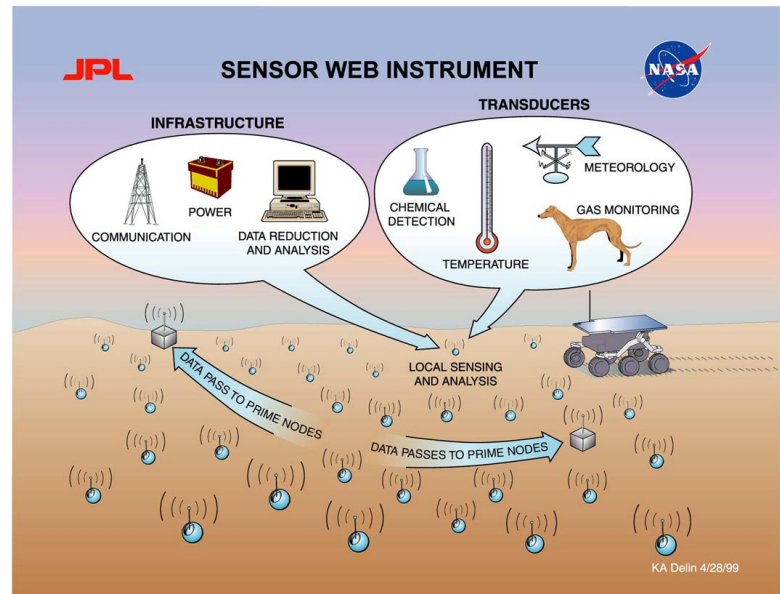


Figure 1: Sensor Web concept as applied to an *in situ* terrestrial environment. The Sojourner rover of the Mars Pathfinder mission is provided for scale.

NASA has occasionally referred to a group of orbiting satellites as a “sensor web”.³ Although this system could be a type of Sensor Web, it is only a subset of the general concept. The connotation of the term “Sensor Web” has a much broader scope and extends from ultra-large-scale *in situ* instruments to correlated remote measurements (such as local spectroscopy) to the coordinated flying (swarming) of distributed spacecraft and everything in between. It is thus too restrictive to say that the Sensor Web concept is a design for an Internet for satellites or an *in situ* instrument, since it can be either one.

2.1 Sensor Webs versus Distributed Sensors

Sensor Webs are often confused with projects that fall under names such as “distributed sensors” or “sensor networks”. The most unique feature of the Sensor Web is that information gathered by one pod is *shared and used* by other pods. Distributed sensors networks merely gather data and communicate it to an uplink point. Examples include the seismology networks present in Southern California. The information gathered by a particular pod on such a network typically does not influence the behavior of another pod. Thus we find that distributed sensors collect data while Sensor Webs modify their behavior on the basis of the collected data. There is a global, macroscopic “purpose” to data collection by Sensor Web pods that is not apparent in the distributed sensor network.

The connectivity in a distributed sensor network is not as integral with its function as it is in the Sensor Web. Unlike those in a sensor network, the individual pods in a Sensor Web matter *to each other*. If, for example, one pod should cease functioning, its lost presence could cause neighboring pods to increase their sampling rate to gain in time resolution what has been lost in spatial resolution. As a second example, consider a Sensor Web whose pods have some limited mobility. By sharing information, the pods can position and reposition themselves along the gradient lines of the environmental parameter

of interest. In this way, the Sensor Web is self-positioning, providing the densest set of mesh points where the parameter of interest is changing the most. This type of global instrument reaction is not possible from a mere set of distributed sensors.

2.2 Multi-Hopping and Power Efficiency

Hopping the data from pod to pod not only allows for sharing locally collected data to other parts of the web but also is power efficient. From elementary electromagnetic theory, the total power required to transmit a signal that ensures a received power P_{rec} a distance D away is $P_{tran} \propto (D/\lambda)^m P_{rec}$, where λ is the wavelength of the transmitted signal.⁴ Here m is 2 in free space and can range to 4 or more in environments with multiple-path interferences or local noise. As a result, the total power required to transmit a given distance with N hops is reduced by a factor of $N^{(m-1)}$ compared to the total power required by direct transmission. In the simple case of free space ($m=2$), this fact is easily demonstrated graphically as shown in Figure 2. It is clear from the figure that the surface area of the collection of smaller spheres (multi-hop transmission) is much less than that of the larger one (direct transmission). The reason why hopping data is power efficient is that more of the power is directed along the path to the receiver. Since the free space case is the best case scenario in terms of transmission efficiency, the value of hopping is only increased in more hostile environments.

FRIS TRANSMISSION EQUATION: $P_{transmit} \propto r^m P_{receive} \quad (2 \leq m \leq 4)$

$$\Rightarrow P_{transmit} \propto \frac{1}{N^{(m-1)}} D^m P_{receive}$$




Figure 2: To communicate data from pod T to pod R, multi-hopping is more power efficient than simple direct transmission. The spheres represent where power is transmitted in the case of free space.

2.3 Sensor Web Organization

The very essence of a Sensor Web, with its multiply redundant pods, allows for an instrument to be “reseeded” against instrument degradation. In other words, as various pods drop out of the web because of spent batteries or damaged transducers, it is possible to deploy new pods in the instrument area that will seamlessly mesh with the existing, older web communication backbone. In this way, though the pods themselves are expendable, the Sensor Web instrument can continue to function indefinitely. Moreover, this reseeding allows the macroinstrument to evolve over time as new pods can be more sophisticated and technologically advanced than older ones. Because the Sensor Web has no definitive boundaries (the mother pods may be located anywhere in the network), multiple deployments of webs in a given area will naturally mesh with one another. Consequently, a Sensor Web represents an instrument whose surveying area can be expanded via multiple deployments.

There is no preferred direction of information flow on the Sensor Web. Although the mother pods provide a point of access into and out of the Sensor Web, there is no assumed directionality or “focusing” the data collected by individual pods towards the mother pods. This allows the Sensor Web instrument the complete flexibility that we have just described: Pods can be dropped in or out at random points in space and time and mother pods may be added to grow the extent of the web. Unlike a star-network, in which a mother pod must be centrally located, mother pods in a Sensor Web may be distributed wherever convenient. The




Figure 3: The scale-free nature of the Sensor Web. Notice the first tier of mother pods (cubes) form their own web structure, leading to the idea that the Sensor Web is a web of web nodes. Individual (spherical) pods are not necessarily associated with a particular (cubic) mother pod, although most tend to be.

assumed omni-directional wireless communication between pods further reduces rigid placement requirements in a Sensor Web.

To formalize these intuitive notions, we define a “pod” as a *physical* piece of hardware containing the two modules of transducers and infrastructure as described above. In contrast, a “node” is a *logical* (or abstract) construct that simply is a vertex on a Sensor Web. Apparently, a node can be either a pod or another Sensor Web. Thus a Sensor Web can be viewed recursively as a web of webs as illustrated in Figure 3. Any “spherical” pods dropped at random into this depicted web will, with a higher probability, tend to be associated with a single mother pod and particular sub-web. Consequently, a scale-free hierarchy emerges which is quite robust with respect to random pod dropout.^{5,6} Interestingly, there are indications that nervous systems may be modeled on similar ideas, again invoking a strong analogy between the Sensor Web and the brain.

3. SENSOR WEBS APPLICATIONS

The Sensor Web concept is not a “one-size-fits-all” solution any more than one would expect all satellites to be the same. It seems natural, for example, to shrink all pods to the smallest size possible. However, in deep ocean applications one wants large pods for logistical reasons. Pod-to-pod baud rates can vary from very low (sonar applications) to the very high (satellite applications). Data transmission schemes in covert military environments are quite different from those on the surface of Mars where there are no FCC regulations. The Sensor Web concept extends from large-scale *in situ* instruments to the coordinated flying (swarming) of a distributed spacecraft. It allows for information synthesis over a large spatial area involving multiple measurements and tracking of dynamic phenomena over multiple length scales simultaneously. The unifying feature of all Sensor Webs, however, is that the information gathered by the instrument is also *used* by the instrument and will modify the behavior of the instrument. The ultimate Sensor Web has an environmental self-awareness and reacts accordingly. A detailed discussion of specific Sensor Web application issues can be found elsewhere.⁷

A Sensor Web provides a different type of measurement tool than that associated with remote measurements made by orbital platforms. By definition, remote measurements obtained from orbit require a high degree of knowledge of the physics of the measurement to infer value from the data collected (interpreting ocean currents or a vegetation index, for example). In contrast, a Sensor Web can provide direct, proximity measurements over a large spatial scale whose value is much more immediate. Moreover, unlike remote measurements made by orbital or airborne platforms, a Sensor Web provides a continued, virtual presence in an area. This is particularly important when investigating phenomena of a transient nature where there is no guarantee that an orbiting instrument will be in position to record the event. Finally, Sensor Webs can provide crucial ground truth and calibration data for remote measurements.

As a specific (though exotic) example of the interplay between a Sensor Web and orbital measurements, consider the problem of searching for life on Mars. Because bacteria are thought to be the most primitive and ancient life forms on Earth, they are the logical target for investigating life via remote sensing. An example of such an instrument for this purpose is the interferometer-based Terrestrial Planet Finder.⁸ This instrument probes atmospheres of extra-solar planets for gases that are out of chemical equilibrium; the disequilibrium potentially caused by biological activity. Planets that harbor significant amounts of bacterial life will possibly have a metabolic signature in their atmospheres. To investigate biological signatures in the Martian atmosphere, however, remote sensing technologies are not adequate because the small volumes of any biologically produced gas disperses rapidly over the surface. Moreover, the likely transient nature of the respired gases might be missed by any orbiting instruments that fly over a specific position only at certain times. Therefore, *in situ* gas sensors are required for detecting traces of any extant biogenic gas emissions. The likely starting place to deploy these sensors would be where water is thought to have once existed. The specific placement is problematic, however, since orbital topographical measurements continue to indicate that there are *vast regions* of Mars that might once have had surface water.⁹ It is simply not practical to land expensive, dedicated spacecraft all over the entire Martian surface. In contrast, Sensor Webs equipped with gas sensors might be able to narrow the search space from those given by the orbital measurements to a much smaller region suitable for a single lander. As a result, the Sensor Web is an instrument capable of scientific surveying on spatial scales between those of an orbital platform and those of a lander.¹⁰

NASA is greatly interested in the Sensor Web concept since it enables a new paradigm for habitat and spacecraft health monitoring and planetary exploration that can significantly impact spacecraft design and mission planning. For example, the Sensor Web addresses NASA’s Human Exploration and Development of Space (HEDS) focus by providing habitat monitoring capabilities (both for crew and hydroponic units) and spacecraft monitoring for diagnostic and warning purposes. Consider, for example, the situation when the Mir Space Station’s hull was breached when it was hit by a cargo vehicle. It was clear to the occupants of the station that air was leaking out, the problem was which module to seal off.¹¹ A Sensor Web

in the station would have greatly aided in leak *location* rather than mere leak *detection*. The Sensor Web concept also specifically addresses NASA's Earth Science Enterprise (ESE) focus of better environmental monitoring and trending and NASA's Space Science Enterprise (SSE) goal of establishing a virtual presence in the solar system. The potential promise of a planetary-scale instrument, as a web of webs, would further solar system exploration immensely. Finally, a Sensor Web can provide a macrointelligence for a robotic craft for applications where "joysticking" is not possible due to fast, dynamic conditions in uncertain environments (such as guiding a rover on a distant planet or landing a spacecraft on an asteroid).

4. SENSOR WEBS AT THE JET PROPULSION LABORATORY

The Jet Propulsion Laboratory (JPL) Sensor Webs Project¹² was formed to specifically address these NASA interests. The goal was to explore the value of the Sensor Web concept in various applications and push the technology of the instrument itself. The JPL Sensor Webs Project would fully leverage the technology revolutions in the computation and telecommunication industries and primarily focus on applying (rather than developing) the required hardware. The applications themselves would determine the specific structure of the pods and the sensors needed.

4.1 Sensor Web 1: A Laboratory Test

Sensor Web 1 was constructed to provide an initial test of simple Sensor Web concepts. A sample pod from this web is shown in Figure 4. Contained in this small package are the transmit/receive chips, a microcontroller, a 3 V Li battery and two sensors to measure the local light level and temperature. The total pod mass is about 50 g. Over a duty cycle of one set of measurements per second, it is estimated that 50 microwatts of power are needed. Further details of this pod may be found elsewhere.⁷

Three identical sensor pods were assembled. A fourth pod was attached to a laptop computer; this served as the prime, or mother, pod where the data were displayed. The protocols developed organized the pod broadcasts by specific time slots. The pod communication range (optimally at 40 m in open field conditions) was intentionally attenuated to allow for a convenient room-sized demonstration.

The 4-node Sensor Web was demonstrated in the two arrangements shown in Figure 5. The first consisted of a linear layout, where each pod could only communicate with its nearest neighbor. The mother pod was able to receive the temperature and light data from all points, thus demonstrating triple hopping as the data was transmitted along the line in a daisy-chain manner. The second demonstration consisted of a dense diamond layout. Here, two of the pods (1 and 2) could communicate with all nodes simultaneously. Because of this redundancy, the removal of either these pods from the web did not interrupt the flow of data from the farthest point on the web (pod 3) to the mother pod, demonstrating the fault tolerant property of the Sensor Web. Moreover, the redundancy in the web creates redundancy in the data transmitted in the web. For example, data from pod 3 can travel to the mother via the route 3-1-M or the longer, and redundant path, 3-2-1-M. The web protocols are designed, however, to eliminate this redundant data dynamically, lest the Sensor Web be overwhelmed with information. Again, the temperature and light levels at all three node points were displayed correctly at the mother pod. In both demonstrated geometries, the mother pod is not located centrally, showing how a Sensor Web is more flexible than a simple star-network.



Figure 4: Functional Sensor Web 1 pod (on left). Extra-terrestrial on right. Note size of quarter.

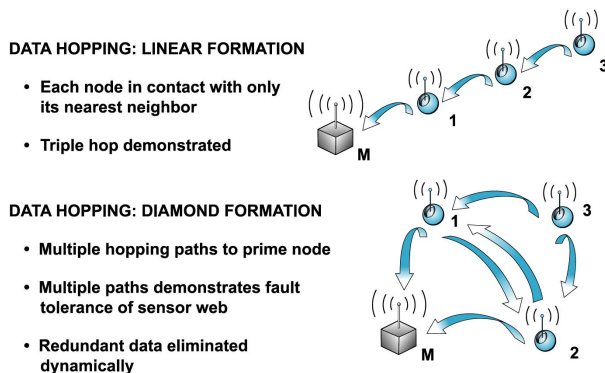


Figure 5: Sensor Web 1 geometries demonstrated. Arrow lengths indicate which pods are in communication range.

4.2 Sensor Web 2: A Botanical Field Test

The success of Sensor Web 1 allowed us to develop a larger system to be deployed in a more realistic application environment. A botanical greenhouse was chosen as the deployment site for several reasons. First, many basic applications, particularly those in astrobiology and earth science, would require a suite of sensors similar to those needed in this environment. Second, a greenhouse would provide somewhat harsh conditions (high temperatures, high humidity, exposure to dust, etc.) to make a field test meaningful while at the same time provide a reasonably protected area for the Sensor Web to operate. Finally, such an environment could be found close to JPL at the Huntington Botanical Gardens in San Marino, California,¹³ allowing for convenient access over the course of the testing.

Prototype pods (see Figure 6) were constructed. A total of seven transducers were used in the initial suite. These commercially available sensors measured relative air humidity, air temperature, soil temperature, soil moisture, local light levels, O₂ gas levels, and H₂S gas levels. The gas sensors were chosen with the idea of ultimately measuring both aerobic and anaerobic activity in some astrobiological applications. All sensors were micropowered with the exception of the H₂S sensor which consumes nearly 1 W because of its hot electrode sensing element. Communications were handled by a pair of commercial receiver and transmitter boards (916 MHz). These micropowered units can handle data at rates of up to 50 kbs and have an open field range of over 150 m. Our application uses burst transmissions at 28.8 kbs. A commercial microcontroller coordinates the radio and sensors and executes the Sensor Web protocols.

The unit is powered by a nominally 8 V battery that is trickle charged by solar cells. When flipped upside down, air-tight food containers were found to provide cheap, light-weight pod packaging. The clear bottom of the containers (which became the top of the Sensor Web pod) were especially convenient as they allowed the solar cells to be protected inside the package while still being fully exposed to the sunlight. White paint was then applied everywhere else to prevent the sunlight from penetrating the clear plastic and directly exposing the electronics inside. Although this packaging would obviously not survive extreme environments for long periods of time, it was adequate for this task and allowed for rapid prototyping. No attempt was made to reduce pod size (surface mounted components were not used, for example) and yet the final Sensor Web 2 pods were relatively small.

To prevent extreme battery drain, each pod monitored its own power consumption and was instructed to sleep should voltage levels drop too much. The pod would wake up after sunlight charged its battery to sufficient levels again. To ensure that the pods would have sufficient power for a relatively long deployment, it was necessary to heat the electrodes in the H₂S sensor only when a measurement was being made. Tests demonstrated that on/off switching of the hot element produced no degradation of sensor performance compared with continuous heating. Despite the success of these tests, the H₂S sensor was eventually




Figure 6: Botanical (Sensor Web 2) pods. The O₂ and H₂S gas sensors are on the left side of the pod while the rechargeable battery is on the right side. The soil temperature thermocouple snakes from the bottom of the package. The small light detector protrudes the top of the pod. Two solar cells charge the battery. Total package dimensions are approximately 5 cm × 10 cm × 16 cm.

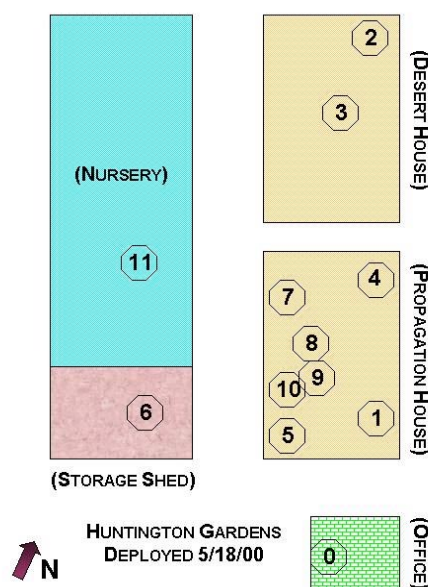


Figure 7: Map of the 12-Pod Sensor Web 2 Layout

dropped from the final pod design because it yielded no information about the specific greenhouse environment at hand.

Sensor Web 2 consisted of 12 pods, including a mother pod that was connected to a laptop computer for data storage and display. It was deployed in the nursery area of the Huntington Gardens and covered an area of approximately 50×100 m² (see Figure 7). The size of this Sensor Web was determined by the application environment and not from limitations of the instrument itself. The application area was chosen because there are four, well-delineated, sub-environments within it. First, a rooftop on a storage shed was available for local weather conditions. Next a four-zone wet (propagation) greenhouse would hold the densest set of pods. The propagation house is a closed environment and misters operated both from the ceiling and ground-level tables creating high relative humidity and warm temperatures. Figure 8 shows pod 8 deployed in this greenhouse. It is situated underneath the sole temperature and relative humidity sensors that monitor the *entire* greenhouse. This pod provides the calibration test of the Sensor Web to the existing greenhouse sensors (which are assumed to be accurate). The third sub-environment is a desert (cactus) greenhouse. Unlike the wet greenhouse, this area was open as the sidewalls were a wire mesh allowing for ambient conditions to exist across the entire greenhouse. For this reason, only two pods were deployed there. Pod 3 was placed under the sole temperature and relative humidity sensors monitoring this area. Finally, the last sub-environment was a moist carnivorous plant bog in an outdoor area covered by a tarp. The mother pod (indicated by a “0” on the map) was located in an office area where electrical power was available for the laptop computer.



Figure 8: Sensor Web 2 pod deployed in a greenhouse. The cylindrical unit hanging from the ceiling is the present sensing system and measures temperature and relative humidity at this one point only for the entire greenhouse.

Sensor Web 2 was activated on May 18, 2000. Every five minutes, 24 hours a day, 7 days a week, each pod recorded the state of its transducers (air temperature, soil temperature, relative humidity, O₂ level, soil moisture, light level) and its internal voltage. This information was then hopped about the Sensor Web until it was captured by the mother pod and recorded. Under typical conditions, a total of 3 hops was required to move data from the farthest point on the web (pod 2) to the mother. The wire-mesh in the desert greenhouse shielded radio communication through the sidewalls and information had to travel through the Plexiglas-like walls on the north and south sides. Although initially planned as a two-week test, Sensor Web 2 proved quite robust and collected data continuously in 5 minute intervals *for over 22 weeks*. The calibration tests at pods 3 and 8 checked out with the existing sensors to within the accuracy of the measurements ($\pm 0.75^{\circ}\text{C}$ for temperature and $\pm 2\%$ for relative humidity).

A number of key concepts were demonstrated by Sensor Web 2:

- 1) Individual pods reintegrate themselves into the web after having been moved around and having been powered off and on. Fault tolerance was demonstrated via multiple communication paths to the mother pod because of the local density of pods.
- 2) Minor hardware repairs, recalibrations, and upgrades have been performed on some of the pods while leaving the rest of the web up and running, and the upgraded pods freely reintegrate with the remainder of the web upon return. This demonstrates the idea that the web itself is the instrument, not the individual pods.
- 3) When pods are powered on but isolated, they seek the rest of the network until such time as they detect neighboring pods. They then integrate themselves into the network. This was accomplished by powering off pods 11, 7, 8, and 4; effectively cutting the web in half. Pods 2 and 3 were isolated. They attempted to contact the rest of the web for several days until their power levels dropped and they went into sleep mode. Pods 11, 7, 8, and 4 were then reactivated. Once charged, pods 2 and 3 reintegrated themselves into the web.

4) Bi-directional flow of information was demonstrated by looking at the data gathered throughout the web via the mother pod, while each of the individual pods were synched to clock pulses emanating from the mother pod. (Recall there is no “data focusing” or preferred direction of information flow on a Sensor Web.) Although each of the pods had its own clock, the mother periodically resynchronized all clocks to itself. This feat was accomplished despite the fact that the synching pulse had to multi-hop out to the various pods. The periodic synchronization reduced the internal clock jitter within the system to less than a few milliseconds as required by the communication protocols. In addition, it guaranteed that the entire collection of pods took their measurements essentially simultaneously (on the order of a millisecond), providing a true large-scale spatial snapshot of the environment.

Sensor Web 2 also has a rudimentary ability to understand itself. Local transmission and reception information is gathered by each of the pods and hopped around the web. Thus the mother pod receives information about the specific path that each piece of sensory data had in traveling through the web. This form of organization is focused on a central point (mother pod). As a continuing part of our program, we currently working on *using* the path information at *every* pod, allowing for the web-level diagnostics to be decentralized and improving communication efficiency.

Typical data from the light sensors on the web is shown in Figure 9. The jagged effect in the data is a result of the shadows cast on the light sensor by the greenhouse superstructure or local plants. Notice how the peaks from pods 1 and 4 (on the east side of the greenhouse) occur in the morning while the peaks in light levels for pod 5 (west side) occur in the afternoon. Thus (not surprisingly) the sun rises in the east and sets in the west. A simple statement but also an example of how the Sensor Web can ultimately provide knowledge of its environment (sun position) rather than mere data (light measurements). Finally note that although pod 11 is more west than pod 5, the effect of the cloth tarp over the outdoor nursery lowers the peak levels and smears things out. To be able to differentiate that environmental difference autonomously presents a well-defined but substantially more difficult problem for future Sensor Web work.




Figure 9: Typical local light level data across Sensor Web 2 as a function of time. Units are relative.

5. CONCLUSION

The Sensor Web represents a new class of instrument that has emerged from the hardware revolutions of the past decade. The JPL Sensor Webs Project has successfully built instruments and demonstrated basic Sensor Web concepts in the field.

These prototype instruments have shown the robustness and potential capabilities of a Sensor Web. At present we are looking towards using the instrument as part of a larger application program. It is expected that the Sensor Web will become a ubiquitous instrument in the future, particular in applications that require an intelligent, virtual presence.

ACKNOWLEDGEMENTS

We gratefully acknowledge James Folsom and Theresa Trunnelle of the Huntington Botanical Gardens for the use of their facilities during the field test and their active and enthusiastic support throughout. We also thank our colleagues at JPL and Caltech who have stimulated various aspects of this work: Greg Bearman, Yi Chao, Pan Conrad, Lea Cox, Jim Edberg, Jason Feldman, Margie Homer, Tom Jedrey, Don Kurtz, Ken Manatt, Gene McDonald, Ken Nealson, Amy Ryan, Lisa Stein, Michael Storrie-Lombardi, Henry Sun, Sasha Tsapin, Rob Stirbl, Bob Van Steenburg, Alice Wessen, and Yuk Yung. We acknowledge continued support from Rod Zieger and Diane Evans in the JPL Earth Science Office. One of us (KAD) is particularly indebted to Arthur (Lonne) Lane, Loren Lemmerman, and Philip Moynihan for many invaluable discussions that have shaped the Sensor Web concept and its potential applications.

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

1. K.A. Delin, S.P. Jackson, and R.R. Some, "Sensor Webs," *NASA Tech Briefs*, **23**, pg. 80, 1999. (See also <http://www.nasatech.com/Briefs/Oct99/NPO20616.html>)
2. C. Koch and G. Laurent, "Complexity and the Nervous System," *Science*, **284**, pp. 96-98, 1999.
3. See, for example, <ftp://ftp.hq.nasa.gov/pub/pao/pressrel/2000/00-065.txt>
4. M.N.O. Sadiku, "Wave Propagation," in *The Electrical Engineering Handbook*, R.C. Dorf, ed., CRC Press, Boca Raton, pp. 925-937, 1997.
5. A-L. Barabasi, R. Albert, and H. Jeong, "Mean-Field Theory of Scale-Free Random Networks," *Physica A*, **272**, pp. 173-187, 1999.
6. R. Albert, H. Jeong, and A-L. Barabasi, "Error and Attack Tolerance of Complex Networks," *Nature*, **406**, pp. 378-382, 2000.
7. K.A. Delin and S.P. Jackson, "Sensor Web for *In Situ* Exploration of Gaseous Biosignatures," *Proceedings of 2000 IEEE Aerospace Conference*, Big Sky, MT March 2000. (See also <http://sensorwebs.jpl.nasa.gov/resources/Delin-Jackson2000.pdf>)
8. C.A. Beichman, N.J. Woolf, and C.A. Lindensmith (eds.), *The Terrestrial Planet Finder*, JPL Publication 99-3, 1999. (See also <http://tpf.jpl.nasa.gov/>)
9. M.C. Malin and K.S. Edgett, "Sedimentary Rocks of Early Mars," *Science*, **290**, pp. 1927-1937, 2000.
10. Ken Nealson, private communication.
11. J. Oberg, "Shuttle-Mir's Lessons for the International Space Station," *IEEE Spectrum*, **35**, pp. 28-37, 1998.
12. For a full description, see <http://sensorwebs.jpl.nasa.gov/>
13. See <http://www.huntington.org/BotanicalDiv/HEHBotanicalHome.html>